

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
<small>Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.</small>				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE 2 Dec 1996	3. REPORT TYPE AND DATES COVERED final: 1 June 1990 - 15 June 1995	
4. TITLE AND SUBTITLE Final Report Automated Model Management / AASERT Supplement			5. FUNDING NUMBERS G - N00014-90-J-1904	
6. AUTHOR(S) Daniel Sabey Weld				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) University of Washington Department of Computer Science & Engineering Box 352350, 114 Sieg Hall Seattle, WA 98195-2350			8. PERFORMING ORGANIZATION REPORT NUMBER G - N00014-90-J-1904	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Dept. of the Navy Office of the Chief of Naval Research 800 North Quincy Stree, Code 1511:KIJ Arlington, VA 22217-5000			10. SPONSORING / MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES				
19961209 046				
a. DISTRIBUTION / AVAILABILITY STATEMENT no limitations			12. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) The objective of our research was to investigate and automate the cognitive processes underlying engineering & manufacturing activities. We sought also to develop "spinoff" technologies that enable advanced engineering tools. In this vein, we focussed on the planning process while continuing our efforts on simulation compilation and automated design. In the latter stages of the grant we shifted emphasis towards the software agent application area.				
DTIC QUALITY INSPECTED 3				
14. SUBJECT TERMS simulation compilation automated design software agent application			15. NUMBER OF PAGES 15	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT	

PI Name: Daniel S. Weld
PI Institution: University of Washington
PI Phone Number: (206) 543-9196
PI E-mail Address: weld@cs.washington.edu
Grant Title: Automated Model Management / ASSERT Supplement
Grant Number: N00014-90-J-1904&P00001
Reporting Period: Final Report (6/1/90-6/15/95)

1 Numerical Productivity Measures

Refereed papers submitted but not yet published: 0
Refereed papers published: 33
Unrefereed reports and articles: 3
Books or parts thereof submitted but not yet published: 0
Books or parts thereof published: 1
Patents filed but not yet granted: 0
Patents granted (include software copyrights): 0
Invited presentations: 23
Contributed presentations: 8
Honors received: 4

- Program Chair, AAAI-96
- Councilor, AAAI-96
- Associate Editor, *Journal of Artificial Intelligence Research*
- Program Chair, QR-93

Prizes or awards received: 0
Promotions obtained: 1

- Assistant Professor to Associate Professor with Tenure.

Graduate students supported: 9
Post-docs supported: 0
Minorities supported: 1

PI Name: Daniel S. Weld
PI Institution: University of Washington
PI Phone Number: (206) 543-9196
PI E-mail Address: weld@cs.washington.edu
Grant Title: Automated Model Management / ASSERT Supplement
Grant Number: N00014-90-J-1904&P00001
Reporting Period: Final Report (6/1/90-6/15/95)

2 Detailed Summary of Technical Results

The objective of our research was to investigate and automate the cognitive processes underlying engineering and manufacturing activities. We sought also to develop "spinoff" technologies that enable advanced engineering tools. In this vein, we focussed our attention on the planning process while continuing our efforts on simulation compilation and automated design. In the latter stages of the grant we shifted emphasis towards the software agent application area.

2.1 Planning with Continuous Change

With J. Scott Penberthy, I designed a partial order planner, ZENO, that efficiently handles domain axioms and actions with complex, conditional and metric effects that take place over extended intervals of time. ZENO reasons about simultaneous actions including continuous change, external events, and actions that interfere with each other. While other planners exist with combinations of these features, ZENO is different because it is arguably sound and complete, yet reasonably efficient.

To achieve this performance, ZENO incorporates three novel features:

1. a technique called *token reduction* matches equational preconditions to relevant effects,
2. fast, incremental algorithms from operations research and logic programming solve metric constraints, and
3. lazy evaluation combined with the Mean Value Theorem deal with continuous change and interval constraints.

We have proven that the ZENO algorithm is sound and complete (modulo a few restrictions on the action representation). In addition, ZENO has been completely implemented in Common Lisp and been tested in numerous domains. Our techniques allow it to process plans requiring only simple change in roughly the same time as required by the UCPOP planner; hence, the cost of reasoning about continuous time is only paid when strictly necessary.

2.2 Probabilistic Planning

Classical planning assumes complete and deterministic information about the world state and the effects of actions. These assumptions are inappropriate for

many domains: turning the ignition key might usually start one's old car, but occasionally fail for unknown reasons. Even if a deterministic model is possible for a given domain, it might be too complex to be useful. For example, when deciding between an indoor and an outdoor site for a wedding, one is likely to use a probabilistic model to forecast the weather rather than project the cloud dynamics. The initial world state is also a source of uncertainty: will the freeways be crowded?

With Steve Hanks and Nick Kushmerick, I devised a planning algorithm that doesn't depend on the assumptions of complete and deterministic information. We use a probability distribution over possible worlds to model imperfect information about the initial world state, and we model actions using a conditional probability distribution over changes to the world.

Adopting a probabilistic model complicates the definition of plan success. Instead of terminating when it builds a plan that *provably* achieves the goal, our planner terminates when it builds a plan that is *sufficiently likely* to succeed: our algorithm produces a plan such that the probability of the plan achieving the goal exceeds a user-supplied probability threshold, if such a plan exists. Our work makes several important contributions.

- We have defined a symbolic action representation and provided it with probabilistic semantics.
- We have developed an algorithm, BURIDAN¹, for probabilistic planning and proven the algorithm is both sound and complete.
- We have implemented BURIDAN in Common Lisp and analyzed its performance. In particular, we've compared the efficiency of four different probabilistic assessment algorithms both analytically and with empirical experiments. In addition, we have explored the interface between the process of generating plans and the process of evaluating them.

2.3 Formal Foundations of Case-Based Planning

With Steve Hanks I developed a formal foundation for the field of case-based planning. Case-based planning is an attractive reasoning paradigm for several reasons. First, cognitive studies have shown that human experts depend on a knowledge of past cases for good problem-solving performance. Second, computational complexity arguments show that reasoning from first principles requires time exponential in the size of the problem. Case-based reasoning systems may avoid this problem by solving a smaller problem: that of adapting a previous solution to the current task. Intuition tells us that most new problem-solving situations closely resemble old situations, therefore there is a clear advantage to using past successes to solve new problems. However, while research efforts in case-based reasoning have made advances in various problem domains, general principles and domain-independent algorithms have been slower to emerge. We explore the theoretical foundations of case-based planning, in particular characterizing the fundamental tradeoffs that govern the process of plan adaptation.

¹Jean Buridan (bō rē dān'), 1300-58, a French philosopher and logician, has been credited with originating probability theory. He seems to have toyed with the idea of using his theory to decide among alternative courses of action: he is credited with the story of "Buridan's Ass," in which an ass that lacked the ability to choose starved to death when placed between two equidistant piles of hay.

Our main result is a domain-independent algorithm, SPA², for plan adaptation that is sound, complete, and systematic. Soundness means that the output plan is guaranteed to satisfy the goal, completeness means that the planner will always find a plan if one exists (regardless of the library plan provided by the retriever), and systematicity means that the algorithm explores the space of adaptations non-redundantly (in short, it will never consider an adaptation more than once).

Systematicity is the most tricky property to guarantee, and for two reasons: first, the adapter operates in a space of incomplete plans³. Each incomplete plan can expand into an exponential number of completions; systematicity requires that the adaptation algorithm never consider two incomplete plans that share even one completion, while completeness requires that every potential completion be considered. Second, plan adaptation requires a combination of retracting previous planning decisions (choice and ordering of plan steps, binding of variables within the operator schemas), as well as making new decisions. Systematicity requires that a decision, once retracted, never be considered again. Achieving this property without incurring a tremendous amount of bookkeeping overhead is a triumph of our algorithm.

We implemented our algorithm in Common Lisp and tested it in both the blocks world and a more sophisticated transportation scheduling domain. Experimental studies compare our algorithm to a similar effort by Kambhampati and Hendler. Our results show a systematic speedup from plan reuse for certain simple and regular problem classes. The simplicity and formal properties of the SPA algorithm make it an ideal framework for characterizing the behavior of existing case-based planners. For example, we analyze Hammond's CHEF system, showing the relationship between its plan-repair strategies and SPA's simpler search algorithm, and demonstrating that CHEF's repair strategies are incomplete.

2.4 Tractable Reasoning about Locally Complete Information

Agents that utilize classical planners presuppose correct and complete information about the world. Having complete information facilitates planning since the agent need not obtain information from the external world — all relevant information is present in the agent's world model (this is the infamous *closed-world assumption* [10]).

Unfortunately, in many important cases, an agent does not have complete information about its world. For instance, a robot may not know the size of a bolt or the location of an essential tool [15, 11]. Similarly, a software agent, such as the UNIX *softbot* (software robot) [9, 8], cannot be familiar with the contents of all the bulletin boards, FTP sites, directories, and files accessible through the Internet.

Recent work has sketched a number of algorithms for planning with incomplete information (*e.g.*, [3, 15, 11, 17, 6, 8]). Unfortunately, these planners are vulnerable to *redundant information gathering* when they plan to "sense" information that is already present in the agent's world model [7, 12, 13]. Since satisfying

²SPA stands for "systematic plan adaptor."

³An incomplete plan may be partially ordered, may contain partially constrained variables, and may require additional steps or constraints for it to achieve the goal.

the preconditions of an information-gathering action can involve arbitrary planning, the cost of redundant information gathering is unbounded in theory and quite large in practice.

With Oren Etzioni and Keith Golden, I developed and implemented a solution to this problem. Our approach utilizes an explicit database \mathcal{D}_G of meta-level sentences such as “I know all the files in /kr94,” which encode *locally complete information*. The information in this database is similar to “closed roles” found in knowledge-representation systems such as CLASSIC [4], LOOM [5], and others. However, our approach has the following novel features:

- A sound and tractable calculus for answering queries based on the information in \mathcal{D}_G . The calculus answers queries such as: if the agent is familiar with all the files in the directory /kr94, and with all group-readable files on the file system, does it follow that the agent is familiar with all the group-readable files in /kr94? Our calculus is incomplete, but we prove that sound and complete query answering is uncomputable.
- A sound and tractable calculus for updating the \mathcal{D}_G as the state of the world changes. The update calculus answers questions such as: if a file is deleted from /kr94, is the agent still familiar with all the group-readable files in /kr94? What if a file is added to /kr94?
- This machinery is incorporated into an extended version of the UCPOP partial-order planner [16] enabling it to satisfy universally quantified goals and avoid redundant information gathering despite the absence of complete information. We measure the impact of the machinery on the planner, utilizing a suite of test problems in the UNIX domain, and show how the benefit of avoiding redundant information gathering far outweighs the costs associated with maintaining and utilizing the \mathcal{D}_G .

2.5 Model Dimensions and Accuracy

I investigated dimensions characterizing engineering models and devised techniques for automatically selecting models based on their accuracy. The overall framework postulates a series of almost orthogonal dimensions: model scope, range, resolution, accuracy, cost, and ontology.

My most detailed analysis concerns the nature of simplifying assumptions. I used Addanki’s *graph of models* (GoM) framework, a directed graph whose nodes represent models of a device and whose arcs are labeled with the simplifying assumption(s) that are added when traversing the edge between the two models. The GoM is best thought of with the simplest models at the top, the most complex models at the bottom and the edges pointing upwards. In this context, model selection can be seen as a search process through the graph. Model refinement techniques move from the current model downward, while model simplification techniques move upwards. We argue that model simplification should be done in the context of a particular goal and that simplification operators should provide some sort of guarantee on the accuracy of the simplified model. I achieved four principle results:

- The notion of *model sensitivity analysis* (MSA) — the problem of predicting how a change in models will affect the resulting predicted behavior over time.

- The definition of a class of model relations called *fitting approximations* for which the problem of model sensitivity analysis is reduced to one of computing the sign of partial derivatives in a single model.
- An implemented, domain-independent theory of query-directed model simplification that uses *bounding abstractions* to provide a guarantee on the answers produced in the approximate model. Model sensitivity analysis is used to compute the bounding abstractions.
- An implemented, domain-independent theory of discrepancy-driven model refinement which also uses model sensitivity analysis.

2.6 Self-Explanatory Simulators

With Franz Amador, I devised and implemented the PIKA self explanatory simulation system [2, 1]. Our technique utilizes an innovative incremental constraint system for determining the causal order of simultaneous equations. This greatly increases the speed of simulation during shifts in operating regions.

2.7 Innovative Design

With Dorothy Neville, I investigated SIE a program that automates first-principles, design of lumped parameter models [14]. The implemented code has proven to be a useful testbed for empirically evaluating the heuristic adequacy of William's ideas about Interaction Topologies.

PI Name:	Daniel S. Weld
PI Institution:	University of Washington
PI Phone Number:	(206) 543-9196
PI E-mail Address:	weld@cs.washington.edu
Grant Title:	Automated Model Management / ASSERT Supplement
Grant Number:	N00014-90-J-1904&P00001
Reporting Period:	Final Report (6/1/90-6/15/95)

3 Publications, Presentations and Reports

Books and Book Chapters:

Weld, D. and Addanki, S. "Query-Directed Approximation" in *Recent Advances in Qualitative Physics*, 101-116, B. Faltings and P. Struss (editors), MIT Press Inc., Cambridge, MA, 1992.

Weld, D., "Theories of Comparative Analysis," MIT Press, Cambridge, MA, May 1990; 184 pages.

Articles in Top Refereed Journals:

Etzioni, O., Golden, K. and Weld, D., "Sound and Efficient Closed-World Reasoning for Planning," to appear in *Artificial Intelligence*, 36 pages, 1996.

Kushmerick, N., Hanks, S. and Weld, D., "An Algorithm for Probabilistic Planning," *Artificial Intelligence*, 76:239-286, July 1995.

Hanks, S. and Weld, D., "A Domain-Independent Algorithm for Plan Adaptation," *Journal of AI Research*, 2:319-360, January 1995.

Barrett, A. and Weld, D., "Partial Order Planning: Evaluating Possible Efficiency Gains," *Artificial Intelligence*, 67:71-112, May 1994.

Weld, D., "Reasoning about Model Accuracy," *Artificial Intelligence*, 56:255-300, August 1992.

Weld, D., "Qualitative Physics: Albatross or Eagle?" *Computational Intelligence*, 8:175-186, May 1992.

Weld, D., "Exaggeration," *Artificial Intelligence*, 43:311-368, June 1990.

Weld, D., "Comparative Analysis," *Artificial Intelligence*, 36:333-374, October 1988. Reprinted as a chapter in *Readings in Qualitative Reasoning about Physical Systems*, D. Weld and J. de Kleer (editors) Morgan Kaufmann, San Mateo, CA, 1989.

Weld, D., "Choices for Comparative Analysis: DQ Analysis or Exaggeration?" *International Journal for Artificial Intelligence in Engineering*, 3:174-180, July 1988.

Weld, D., "The Use of Aggregation in Causal Simulation," *Artificial Intelligence*, 30:1-34, October 1986. Reprinted as a chapter in *Readings in Qualitative Reasoning about Physical Systems*, D. Weld and J. de Kleer (editors) Morgan Kaufmann, San Mateo, CA, 1989.

Articles in Major Refereed Conferences:

Draper, D., Hanks, S. and Weld, D., "A Probabilistic Model of Action for Least-Commitment Planning with Information Gathering," *Tenth Conference on Uncertainty in Artificial Intelligence (UAI-94)*, pages 178-186, Seattle, WA, July 1994.

Weld, D. and Etzioni, O., "The First Law of Robotics (a call to arms)," *Twelfth National Conference on Artificial Intelligence (AAAI-94)*, pages 1042-1047, Seattle, WA, July 1994.

Barrett, A. and Weld, D., "Task-Decomposition via Plan Parsing," *Twelfth National Conference on Artificial Intelligence (AAAI-94)*, pages 1117-1122, Seattle, WA, July 1994.

Penberthy, J. and Weld, D., "Temporal Planning with Continuous Change," *Twelfth National Conference on Artificial Intelligence (AAAI-94)*, pages 1010-1015, Seattle, WA, July 1994.

Kushmerick, N., Hanks, S. and Weld, D., "An Algorithm for Probabilistic Least-Commitment Planning," *Twelfth National Conference on Artificial Intelligence (AAAI-94)*, pages 1073-1078, Seattle, WA, July 1994.

Golden, K., Etzioni, O. and Weld, D., "Omnipotence Without Omniscience: Efficient Sensor Management for Planning," *Twelfth National Conference on Artificial Intelligence (AAAI-94)*, pages 1048-1054, Seattle, WA, July 1994.

Draper, D., Hanks, S. and Weld, D., "Probabilistic Planning with Information Gathering and Contingent Execution," *Second International Conference on AI Planning Systems (AIPS-94)*, 31-36, Chicago, IL, June 1994.

Etzioni, O., Golden, K. and Weld, D., "Tractable Closed World Reasoning with Updates," *Fourth International Conference on Principles of Knowledge Representation and Reasoning (KR-94)*, 178-189, Bonn, Germany May 1994.

Barrett, A. and Weld, D., "Characterizing Subgoal Interactions for Planning," *Thirteenth International Joint Conference on Artificial Intelligence (IJCAI-93)*, 1388-1393, Chambery, France, August 1993.

Amador, F., Finkelstein, A. and Weld, D., "Real-Time Self Explanatory Simulation," *Eleventh National Conference on Artificial Intelligence (AAAI-93)*, 562-567, Washington, D.C., July 1993.

Neville, D. and Weld, D., "Innovative Design as Systematic Search," *Eleventh National Conference on Artificial Intelligence (AAAI-93)*, 737-742, Washington, D.C., July 1993.

Sun, Y. and Weld, D., "A Framework for Model-Based Repair," *Eleventh National Conference on Artificial Intelligence (AAAI-93)*,

182-187, Washington, D.C., July 1993.

Penberthy, J. S. and Weld, D., "UCPOP: A Sound, Complete, Partial-Order Planner for ADL," *Third International Conference on Principles of Knowledge Representation and Reasoning (KR-92)*, 103-114, Cambridge, MA, October 1992.

Etzioni, O., Hanks, S., Weld, D., Draper, D., Lesh, N. and Williamson, M., "An Approach to Planning with Incomplete Information," *Third International Conference on Principles of Knowledge Representation and Reasoning (KR-92)*, 115-125, Cambridge, MA, October 1992.

Hanks, S. and Weld, D., "Systematic Adaptation for Case-Based Planning," *First International Conference on AI Planning Systems (AIPS-92)*, 96-105, College Park, MD, June 1992.

Weld, D., "Approximation Reformulations," *Eighth National Conference on Artificial Intelligence (AAAI-90)*, 407-412, Boston, Massachusetts, August 1990.

Other Refereed Papers:

Weld, D., "An Introduction to Least Commitment Planning," *AI Magazine*, 15:4, 27-61, Winter 1994.

Amador, F., Berman, D., Borning, A., DeRose, T., Finkelstein, A., Neville, D., Notkin, D., Salesin, D., Salisbury, M., Sherman, J., Sun, Y., Weld, D. and Winkenbach, G., "Electronic 'How Things Work' Articles: A Preliminary Report," *IEEE Transactions on Knowledge and Data Engineering*, 5:600-610, August 1993.

Weld, D. and Penberthy, J., "A New Approach to Temporal Planning (Preliminary Report)," *1993 AAAI Spring Symposium on Foundations of Automatic Planning: The Classical Approach and Beyond*, 112-117, Stanford, CA, 1993.

Weld, D., "Generating Simplified Models with Confidence," *American Society of Mechanical Engineers, Symposium on Automated Modeling (ASME DSC-Vol. 41)*, 21-29, Anaheim, CA, November 1992.

Neville, D. and Weld, D., "Innovative Design as Systematic Search," *1992 AAAI Fall Symposium on Design from Physical Principles*, 19-24, Cambridge, MA, October 1992.

Sun, Y. and Weld, D., "Beyond Simple Observation: Planning to Diagnose," *Third International Workshop on Principles of Diagnosis (DX-92)*, 66-75, Orcas Island, WA, October 1992.

Weld, D., "Automatic Selection of Bounding Abstractions," *Sixth International Workshop on Qualitative Reasoning (QR-92)*, 36-43, Edinburgh, Scotland, August 1992. (Also appears in *AAAI-92 Workshop on Approximation and Abstraction*.)

Hanks, S. and Weld, D., "Systematic Plan Adaptation," *1992 AAAI Spring Symposium on Computational Considerations in Supporting Incremental Modification*, 13-18, Stanford, CA, March 1992.

Amador, F. and Weld, D., "Microscopic and Statistical Reasoning about Populations," *AAAI-91 Workshop on Model-Based Reasoning*, 11-16, Anaheim, CA, July 1991.

Weld, D., "System Dynamics and the Qualification Problem" *International Journal of Expert Systems* 3:425-435, 1990. Reprinted as a chapter in *Reasoning Agents in a Dynamic World: The Frame Problem*, K. Ford and P. Hayes (editors), JAI Press Inc., Greenwich, CT, 1991.

Weld, D. and Addanki, S., "Task-Driven Model Abstraction," *Fourth International Workshop in Qualitative Physics*, pages 16-30, Lugano, Switzerland, July 1990.

Amador, F. and Weld, D., "Multi-level Modeling of Populations," *Fourth International Workshop in Qualitative Physics*, 210-219, Lugano, Switzerland, July 1990. (also appears in the *AAAI-90 Workshop on Model-Based Reasoning*, July 1990.)

Reviews and Opinions:

Etzioni, O. and Weld, D., "A Softbot-Based Interface to the Internet," *Communications of the ACM*, 37(7), 72-76, July 1994. Reprinted as a chapter in *Advanced Planning Technology*, A. Tate (editor), The AAAI Press, Menlo Park, CA., USA, May 1996, ISBN 0-929280-98-9.

Weld, D., "The Seventh International Workshop on Qualitative Physics," *AI Magazine*, 15(2), page 88, Summer 1994.

Weld, D., Book Review of "Representations of Commonsense Reasoning," by Ernest Davis, *Artificial Intelligence*, 61:113-120, 1993.

Weld, D., "Comparative Analysis and Qualitative Reasoning," *IEEE Expert*, pages 5-8, February 1991.

Weld, D., Book Review of "The Psychology of Everyday Things," by Donald Norman, *Artificial Intelligence*, 41:111-114, November 1989. Reprinted as a chapter in *Contemplating Minds: Forum in Artificial Intelligence*, M. Stefik, B. Clancy and S. Smoliar (editors), MIT Press, Cambridge, MA, 1993.

Weld, D., Book Review of "Women, Fire, and Dangerous Things" by George Lakoff, *Artificial Intelligence*, 35:137-141, May 1988. Reprinted as a chapter in *Contemplating Minds: Forum in Artificial Intelligence*, M. Stefik, B. Clancy and S. Smoliar (editors), MIT Press, Cambridge, MA, 1993.

PI Name: Daniel S. Weld
PI Institution: University of Washington
PI Phone Number: (206) 543-9196
PI E-mail Address: weld@cs.washington.edu
Grant Title: Automated Model Management / ASSERT Supplement
Grant Number: N00014-90-J-1904&P00001
Reporting Period: Final Report (6/1/90-6/15/95)

4 Research Transitions and DoD Interactions

I have had close interaction with researchers at Xerox PARC and advise on modeling issues in their automated copier diagnosis project.

I have collaborated with researchers at IBM T.J. Watson Research Center in Hawthorne, NY on issues in planning and design.

I collaborate with researchers at Rockwell International on planning algorithms. With Professors Etzioni and Hanks, I recently received a grant from ARPA / Rome Labs supporting work on software agents.

PI Name: Daniel S. Weld
PI Institution: University of Washington
PI Phone Number: (206) 543-9196
PI E-mail Address: weld@cs.washington.edu
Grant Title: Automated Model Management / ASSERT Supplement
Grant Number: N00014-90-J-1904&P00001
Reporting Period: Final Report (6/1/90-6/15/95)

5 Software and Hardware Prototypes

We developed a new release of our portable Common Lisp implementation of the UCPOP planner. New features in version 4.0 include:

- A sophisticated graphical plan debugger.
- Safety constraints [18].
- Quantification over dynamic universes (object creation and destruction).
- New search control strategies and improved search control language.
- Improved user's manual and more domain theories for testing.
- Domain axioms.
- Procedural attachment (predicates expanding to lisp code).
- Declarative specification of search control rules
- Larger set of domain theories & search functions for testing.
- Improved speed.

Since distribution is via anonymous FTP, we suspect that many people have acquired our software without our knowledge, but we estimate that over 140 sites are using the software. Confirmed users include: Air Force Rome Laboratories, Boeing Advanced technology Center, Rockwell International, Jet Propulsion Laboratory, Xerox Corporation, IBM Corporation, AT&T Bell Labs, Massachusetts Institute of Technology, Stanford, University of California at Berkeley, Carnegie Mellon University, University of California at San Diego, University of Southern California / ISI, University of Texas at Austin, Northwestern University, Arizona State University, Brown University, University of Pittsburgh, University of Minnesota, and researchers in France, Germany, Switzerland, Canada, Denmark, and Korea. No direct commercialization is anticipated, but we licensed the code to Apple Computer.

We have implemented a planner handling incomplete information XII, but no commercialization is anticipated.

We have implemented a local closed world reasoning module; Apple computer has indicated interest in licensing the code.

We have implemented a probabilistic planner, called BURIDAN, but no commercialization is anticipated.

We have implemented a temporal planner, called ZENO, but no commercialization is anticipated.

We have improved our implementation of the SIE program for innovative design. No commercialization is anticipated.

Using a combination of C, Common Lisp, and Mathematica, we have implemented PIKA — a self explanatory simulation compiler. No commercialization is anticipated.

References

- [1] Franz G. Amador, Deborah Berman, Alan Borning, Tony DeRose, Adam Finkelstein, Dorothy Neville, David Notkin, David Salesin, Mike Salisbury, Joe Sherman, Ying Sun, Daniel S. Weld, and Georges Winkenbach. Electronic "how things work" articles: Two early prototypes. *IEEE Transactions on Knowledge and Data Engineering*, 5(4):611–618, August 1993.
- [2] Franz G. Amador, Adam Finkelstein, and Daniel S. Weld. Real-time self-explanatory simulation. In *Proc. 11th Nat. Conf. on AI*, pages 562–567, July 1993. Available via FTP from pub/ai/ at ftp.cs.washington.edu.
- [3] J Ambros-Ingerson and S. Steel. Integrating planning, execution, and monitoring. In *Proc. 7th Nat. Conf. on AI*, pages 735–740, 1988.
- [4] R. Brachman. "Reducing" CLASSIC to Practice: Knowledge Representation Theory Meets Reality. In *Proc. 3rd Int. Conf. on Principles of Knowledge Representation and Reasoning*, October 1992.
- [5] D. Brill. *LOOM Reference Manual*. USC-ISI, 4353 Park Terrace Drive, Westlake Village, CA 91361, version 1.4 edition, August 1991.
- [6] Oren Etzioni, Steve Hanks, Daniel Weld, Denise Draper, Neal Lesh, and Mike Williamson. An approach to planning with incomplete information. In *Proc. 3rd Int. Conf. on Principles of Knowledge Representation and Reasoning*, pages 115–125, 1992.
- [7] Oren Etzioni and Neal Lesh. Planning with incomplete information in the UNIX domain. In *Working Notes of the AAAI Spring Symposium: Foundations of Automatic Planning: The Classical Approach and Beyond*, pages 24–28, Menlo Park, CA, 1993. AAAI Press.
- [8] Oren Etzioni, Neal Lesh, and Richard Segal. Building softbots for UNIX (preliminary report). Technical Report 93-09-01, University of Washington, 1993. Available via anonymous FTP from pub/etzioni/softbots/ at cs.washington.edu.
- [9] Oren Etzioni and Richard Segal. Softbots as testbeds for machine learning. In *Working Notes of the AAAI Spring Symposium on Knowledge Assimilation*, pages 43–50, 1992.
- [10] M. Ginsberg, editor. *Readings in Nonmonotonic Reasoning*. Morgan Kaufmann, San Mateo, CA, 1987.
- [11] K. Krebsbach, D. Olawsky, and M. Gini. An empirical study of sensing and defaulting in planning. In *Proc. 1st Intl. Conf. on AI Planning Systems*, pages 136–144, June 1992.
- [12] M. Mason. Kicking the sensing habit. *AI Magazine*, 14(1):58–59, Spring 1993.
- [13] D. Miller. A twelve-step program to more efficient robotics. *AI Magazine*, 14(1):60–63, Spring 1993.
- [14] D. Neville and D. Weld. Innovative design as systematic search. In *Proc. 11th Nat. Conf. on AI*, pages 737–742. MIT Press(AAAI), July 1993. Available via FTP from pub/ai/ at ftp.cs.washington.edu.

- [15] D. Olawsky and M. Gini. Deferred planning and sensor use. In *Proceedings, DARPA Workshop on Innovative Approaches to Planning, Scheduling, and Control*. Morgan Kaufmann, 1990.
- [16] J.S. Penberthy and D. Weld. UCPOP: A sound, complete, partial order planner for ADL. In *Proc. 3rd Int. Conf. on Principles of Knowledge Representation and Reasoning*, pages 103–114, October 1992. Available via FTP from pub/ai/ at ftp.cs.washington.edu.
- [17] M. Peot and D. Smith. Conditional Nonlinear Planning. In *Proc. 1st Intl. Conf. on AI Planning Systems*, pages 189–197, June 1992.
- [18] Dan Weld and Oren Etzioni. The first law of robotics (a call to arms). In *Proc. 12th Nat. Conf. on AI*, pages 1042–1047, 1994.

FORM A2-2

AUGMENTATION AWARDS FOR SCIENCE & ENGINEERING RESEARCH TRAINING (AASERT)
REPORTING FORM

The Department of Defense (DOD) requires certain information to evaluate the effectiveness of the AASERT program. By accepting this Grant Modification, which bestows the AASERT funds, the Grantee agrees to provide the information requested below to the Government's technical point of contact by each annual anniversary of the AASERT award date.

1. Grantee identification data: (R & T and Grant numbers found on Page 1 of Grant)

- a. University of Washington
University Name
- b. N00014-90-J-1904 Mod.# P00004 c. 400x077vip05
Grant Number R & T Number
- d. Daniel S. Weld e. From: 08/16/93 To: 08/15/94
P.I. Name AASERT Reporting Period

NOTE: Grant to which AASERT award is attached is referred to hereafter as "Parent Agreement."

2. Total funding of the Parent Agreement and the number of full-time equivalent graduate students (FTEGS) supported by the Parent Agreement during the 12-month period prior to the AASERT award date.

- a. Funding: \$ 75,000
- b. Number FTEGS: .52

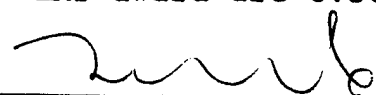
3. Total funding of the Parent Agreement and the number of FTEGS supported by the Parent Agreement during the current 12-month reporting period.

- a. Funding: \$ -0-
- b. Number FTEGS: -0-

4. Total AASERT funding and the number of FTEGS and undergraduate students (UGS) supported by AASERT funds during the current 12-month reporting period.

- a. Funding: \$ 98,677
- b. Number FTEGS: 1.04
- c. Number UGS: 0.0

VERIFICATION STATEMENT: I hereby verify that all students supported by the AASERT award are U.S. citizens.


Principal Investigator

12-2-96
Date

FORM A2-2

AUGMENTATION AWARDS FOR SCIENCE & ENGINEERING RESEARCH TRAINING (AASERT)
REPORTING FORM

The Department of Defense (DOD) requires certain information to evaluate the effectiveness of the AASERT program. By accepting this Grant Modification, which bestows the AASERT funds, the Grantee agrees to provide the information requested below to the Government's technical point of contact by each annual anniversary of the AASERT award date.

1. Grantee identification data: (R & T and Grant numbers found on Page 1 of Grant)

- a. University of Washington
University Name
- b. N00014-90-J-1904 Mod.# P00004 c. 400x077yip05
Grant Number R & T Number
- d. Daniel S. Weld e. From: 08/16/94 To: 08/15/95
P.I. Name AASERT Reporting Period

NOTE: Grant to which AASERT award is attached is referred to hereafter as "Parent Agreement."

2. Total funding of the Parent Agreement and the number of full-time equivalent graduate students (FTEGS) supported by the Parent Agreement during the 12-month period prior to the AASERT award date.

- a. Funding: \$ 75,000
- b. Number FTEGS: .52

3. Total funding of the Parent Agreement and the number of FTEGS supported by the Parent Agreement during the current 12-month reporting period.

- a. Funding: \$ -0-
- b. Number FTEGS: -0-

4. Total AASERT funding and the number of FTEGS and undergraduate students (UGS) supported by AASERT funds during the current 12-month reporting period.

- a. Funding: \$ 98,677
- b. Number FTEGS: .02
- c. Number UGS: 0.0

VERIFICATION STATEMENT: I hereby verify that ~~XXX~~ students supported by the AASERT award are U.S. citizens.

One student under this program was proposed and awarded as a non-U.S. citizen.

[Signature]
Principal Investigator

12-2-90
Date